

APDEC: Algorithms Software

The SciDAC Applied Partial Differential Equations Center for Enabling Technologies (APDEC) strives to advance scientific discovery through the development of applied mathematics. Researchers working on SciDAC Science Application projects and non-SciDAC scientists both depend on APDEC's simulation tools for solving complex problems in scientific fields such as combustion, magnetic fusion, astrophysics, and accelerator modeling.

Applied Math for Applied Science

Many important scientific problems—such as combustion, magnetic fusion, systems biology, and climate change—involve multiple physical processes operating on multiple space and time scales. For these kinds of problems, the components of simulation become quite complicated, as well as strongly linked. The high level of uncertainty and empiricism involved in the design of such mathematical models contrasts the more cooperative “first principles” models used in other areas of science. The choice of model depends not only on the resolved scales, but also on the specific scientific question being asked. Even when the models are well-characterized, there is often an incomplete understanding of fundamental mathematical issues, such as well-posedness.

There is still a great deal of coherence in the underlying mathematical representations of these problems. They are all described in terms of various generalizations of the classical elliptic, parabolic, and hyperbolic partial differential equations (PDE) of mathematical physics. The enormous variety and subtlety found in these applications comes from the way the PDEs are coupled, generalized, and combined with models for other physical processes. The complexity of these models leads to a diverse collection of requirements on the numerical methods, with many open questions about the stability of coupled algorithms. Model complexity also makes it very difficult to obtain high performance. There are tradeoffs between the models, the discretizations,

and the software, with a combination of analysis and computational experiments used to explore the design space. This places a premium on the availability of a diverse and agile software toolset to enable experimentation.

The SciDAC Applied Partial Differential Equations Center for Enabling Technologies (APDEC) is developing a collection of algorithmic and software components that can be assembled to simulate a broad range of complex multicomponent physical systems in which PDEs play a central role. Specifically, APDEC's goal is to enable the agile development of high-performance simulation codes for complex multiphysics and multiscale applications by providing a flexible toolset that meets three requirements—capability, expressiveness, and performance.

Capability involves a sufficient variety of tools to support the broadest possible range of applications, as well as the ability to combine the tools in as many ways as possible. Expressiveness is a concept defined as a factoring of the tools so that nonessential details are hidden from the users who do not need them, while still allowing access to those details for users who want to customize the capabilities in the software. Performance deals with optimizing the implementation of tools to obtain the highest possible performance on high-end computing platforms.

The development of simulation software for complex applications is an end-to-end activity, with the choices regarding models, discretizations, and software all strongly interacting with

APDEC's goal is to enable the agile development of high-performance simulation codes for complex multiphysics and multiscale applications by providing a flexible toolset that meets three requirements—capability, expressiveness, and performance.

and for DISCOVERY




Figure 1. An image from an adaptive mesh refinement (AMR) particle-in-cell (PIC) calculation.

one another, and with fundamental mathematical questions that underlie these choices. For that reason, APDEC researchers work closely with applications developers in all aspects of simulation code design. In particular, members of the APDEC team are taking the lead on the development of new simulation capabilities in two applications areas—magnetohydrodynamic (MHD) modeling of tokamaks, and combustion. In addition, APDEC is collaborating with other SciDAC projects in accelerator modeling and astrophysics to assist them in using various components from the APDEC software suite. Finally, as part of an outreach effort, APDEC is working with investigators on non-SciDAC projects in a broad range

of disciplines to assist in developing new simulation capabilities (sidebar “Outreach Collaborations,” p24).

The APDEC Approach

The APDEC algorithmic approach is based on finite-difference discretizations of PDEs on locally structured grids, combined with techniques for introducing adaptivity and irregular geometry, plus methods for representing kinetic problems using particles.

In grid-based methods for numerical PDEs, it is necessary to make a fundamental choice of discretization technologies. APDEC researchers have chosen to use locally structured grids—ones that

Outreach Collaborations

As part of its outreach activities, APDEC works with scientists from a broad range of disciplines as they attempt to use the Center's software for their particular applications. The following are some of the applications currently under development.

The solution of gyrokinetic equations in four to five dimensions (gyrokinetic phase space) plus time to simulate the dynamics of the plasma in the edge region of a tokamak is part of a collaborative effort with the Edge Simulation Laboratory, a joint project between LBNL and LLNL.

APDEC is working with Dr. Francesco Miniati of the Eidgenössische Technische Hochschule (ETH) in Zürich on AMR cosmology code that couples a finite-volume method for compressible flow with a cell-centered AMR-PIC algorithm for collisionless matter, with the two components coupled through self-gravity.

In collaboration with researchers at Stanford University, APDEC is pursuing simulations of nonlinear propagation of internal waves over long distances in the ocean, using embedded boundary AMR algorithms and software to represent stratified

flows with realistic representations of bathymetry.

A joint project between LBNL and the California Department of Water Resources involves embedded boundary AMR for two-dimensional shallow-water equations in realistic geometries to model the hydrology of the San Francisco Bay and Delta.

APDEC researchers are working together with Dr. Fotini Chow, University of California—Berkeley, to develop simulation models that are insensitive to discontinuous changes in resolved length scales.

APDEC researchers work closely with applications developers in all aspects of simulation code design.

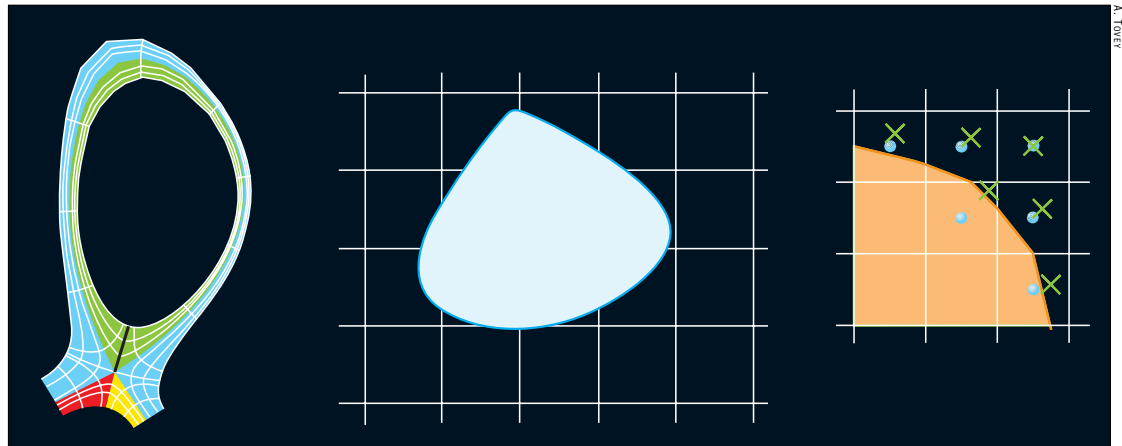


Figure 2. Representations of irregular geometry using structured grids. On the left, a mapped multiblock grid, with grid lines in white, and the different blocks colored green, blue, red, and yellow. In the center, a node-centered cut-cell grid. Unknowns are defined at intersections of grid lines, with special difference approximations applied at the nodes near the irregular boundary. On the right, a volume-of-fluid cut-cell discretization. Exposed portions of cubic control volumes not covered by the body—shown in orange—define the control volumes. Primary dependent variables are defined at Cartesian cell centers (solid circles) and the divergence of fluxes is defined at the centroids (Xs).

are based on defining discrete unknowns on a rectangular discretization of the spatial independent variables. Specifically, APDEC will mostly be using the finite-volume approach, in which the rectangular grid defines a collection of control volumes, for which a natural discretization of the divergence operator is obtained by integrating over the control volume. This leads to methods that satisfy discrete conservation laws, an essential feature for computing discontinuous solutions to PDEs and a desirable property for a much larger class of physical problems. There is a large body of experience in how to construct stable and accurate discretizations to PDEs based on this approach for a broad range of physics problems. Regularity in space leads to regular data layouts, making it easier to optimize for various types of data locality. The regular geometric structure of the spatial dis-

tribution of unknowns leads to efficient iterative solvers for elliptic and parabolic problems. Coupling to particle methods is inexpensive and well-understood. Locating the particles on the grid is trivial, and interpolation between particles and fields has an extensive body of practical experience and mathematical analysis.

Treatment of Irregular Geometries

Two approaches are pursued in the treatment of complex geometries, and both are based on the discretizations on locally rectangular grids (figure 2). One is the use of mapped multiblock grids, in which the computational domain is represented as a union of images of smooth maps from rectangles in abstract coordinate spaces, with boundaries that are aligned with one another. The other approach is the cut-cell approach, in which





Figure 3. Examples of cut-cell grid generation; these are grids generated from a trachea with tracheostomy (upper panel) and a magnetic resonance image (MRI) dataset used to compute flow in a cerebral artery (lower panel).

the irregular domain is represented by its intersection with a rectangular Cartesian grid. While this approach can be applied to nodal-point discretizations, efforts are mostly being concentrated on the volume-of-fluid version of cut cells, in which an irregular boundary is represented on a rectangular grid by intersecting each rectangular cell with the boundary. This leads to natural finite-volume discretizations of the solution to the PDEs in the irregular domain that satisfy discrete conservation laws. In the case where the surface being represented in this way is an irregular domain boundary, these methods are often referred to as Cartesian grid or embedded boundary methods.

The multiblocked mapped grid approach is needed in applications for which anisotropies in the model require the use of a grid that is aligned




Figure 4. A block-structured nested refinement. The white grid on the slicing planes define individual control volumes. The black wire-frame shows the organization of the refined region into blocks.

with respect to those anisotropies to maintain accuracy. Multiple blocks used in cases for which a single rectangular coordinate mapping would be too complicated or singular. This is the case for atmospheric fluid dynamics, problems in subsurface flow modeling, and MHD modeling of tokamaks—machines that generate magnetic fields for confining plasma, which is critical to fusion energy research. One of the principal advantages to the cut-cell approach (figure 3) is the ease with which the numerical grid generation problem, often a major bottleneck in other methods, can be solved. Researchers have developed a tool for embedded boundary grid generation from Initial Graphics Exchange Specification (IGES) descriptions of complex engineering geometries that are fully automatic and take only a modest amount of computer time to generate grids for complex engineering shapes. More recently, the APDEC team has developed equally efficient techniques for generating embedded boundary discretizations information from image data or geographical data specified as the zero set of some function. The latter tools can also be used for grid generation from level-set descriptions of moving fronts.

Adaptive Mesh Refinement

The adaptive mesh methods used by APDEC are based on block-structured adaptive mesh refinement (AMR) algorithms. In this approach, the regions to be refined are organized into rectangular patches of several hundred to several thousand grid cells per patch (figure 4). Thus the high-resolution rectangular grid methods described above can be used to advance the solution in time. Furthermore, the overhead in managing the irregular data is amortized over relatively large amounts of floating point work on regular arrays. For time-dependent problems, refinement is performed in time as well as in space. Each level of spatial refinement uses its own stable time step, with the time steps at a level constrained to be integer multiples of the time steps at all finer levels.

One of the principal advantages to the cut-cell approach is the ease with which the numerical grid generation problem, often a major bottleneck in other methods, can be solved.

Developing APDEC Software

The design of the framework on which APDEC software development is based consists of five major components:

Layer 1—Classes for representing data and computations on unions of rectangles, including a mechanism for managing the distribution of rectangular patches across processors, and an interface to Fortran for obtaining acceptable uniprocessor performance. This is meant to support an underlying coarse-grained model of parallelism based on domain decomposition.

Layer 2—Classes for representing interlevel interactions, such as averaging and interpolation, interpolation of coarse–fine boundary conditions, and managing conservation at coarse–fine boundaries.

Layer 3—Interface classes that implement control structures—such as Berger–Olinger time-stepping, or various iterative methods for solving linear systems—without knowing the details of the data, using a combination of inheritance and class templates. Core solver components for elliptic and

hyperbolic PDEs are developed based on these classes.

Layer 4—Implementation of specific applications or classes of applications using these tools, such as an AMR method for hyperbolic conservation laws or for incompressible flow.

Utility Layer—Support for problem setup, input/output, and visualization, including providing interfaces to existing packages in these areas.




Figure 5. The combined AMR and embedded boundary method used to compute the propagation of a supersonic jet into a vacuum. Each box corresponds to a refined grid patch.

AMR is a mature technology for problems without geometry, with a variety of implementations and applications for various nonlinear combinations of elliptic, parabolic, and hyperbolic PDEs. In particular, the collection of APDEC applications addresses most of the major algorithmic issues in developing adaptive algorithms for the applications described above, such as adaptive multigrid solvers for Poisson’s equation, the representation of non-ideal effects in MHD, and the coupling of particle methods to AMR field solvers. AMR has also been successfully coupled to the mapped-grid and volume-of-fluid methods for treating irregular geometries (figure 5).

Software Design

One of the principal characteristics of the algorithms described is that they are difficult to implement, particularly on parallel computers. The algorithms are more complicated than traditional finite-difference methods, and often the data structures involved are not easily represented in the traditional procedural programming environments used in scientific computing. To manage

this algorithmic complexity, a collection of libraries written in a mixture of Fortran and C++ are used to implement a domain-specific set of abstractions for the combination of algorithms. In this approach, the high-level data abstractions are implemented in C++, while the bulk of the floating point work is performed on rectangular arrays by Fortran routines.

This design approach (sidebar “Developing APDEC Software”) is based on two ideas. The first is that the mathematical structure of the algorithm domain maps naturally into a combination of data structures and operations on those data structures, which can be embodied in C++ classes. The second is that the mathematical structure of the algorithms can be factored into a hierarchy of abstractions, leading to an analogous factorization of the framework into reusable components, also called layers. This reusability is realized by a combination of generic programming and sub-classing. A principal advantage to this design is the relative stability of the application programming interfaces (API) as seen by the applications developer. Implementations may change considerably to enhance performance or in response to changes in the architecture, but these changes are less likely to cause major upheavals to the applications programs. This is because the APIs reflect the mathematical structure of the algorithms, which remain a relatively fixed target.

The design of the APDEC AMR infrastructure has a strong focus on high performance. At the single processor level, most of the floating point work on multidimensional arrays is performed by calls to optimized Fortran. Data holders defined on unions of rectangles are automatically distributed over processors, and the predefined

As with many scientific endeavors that are simply not practical for the laboratory bench, computational experiments have emerged as the only reliable tools for developing verifiable, quantitative predictions.




Figure 6. A time sequence, from left to right, of a vortex ring merger problem using the AMR incompressible Navier–Stokes code included as part of the APDEC code release. There are three grid levels, with a factor of four refinement between successive levels. Middle level boxes are shown in green, finest level boxes in blue, and a single vorticity isosurface in yellow.

aggregate copy operations include the communications calls implemented using Message Passing Interface (MPI). These features, combined with tools for load balancing, lead to software whose computational cost per grid point and scalability is comparable to that of uniform-grid calculations using the same algorithms. For the original Berger–Colella algorithm for hyperbolic conservation laws, 75% efficiency was measured on up to 1,024 processors on IBM SP and Hewlett-Packard Alpha systems, and up to 32,000 processors on the IBM BlueGene/L using specialized communication primitives available on that system. For low-Mach number combustion, production calculations on IBM SP and Silicon Graphics (SGI) systems are typically done using between 512 and 4,096 processors.

All of the APDEC software resides in a Concurrent Versions System (CVS) repository, with checkout access to collaborators through remote CVS access. As various components of the software reach the appropriate state of hardening, they are made publicly available under a FreeBSD license from the APDEC website. The current public release version contains implementations of the core Layer 1 through Layer 3 libraries, including solvers for elliptic and hyperbolic PDEs on unions of rectangles and AMR grid hierarchies. As part of the distribution APDEC also provides a number of complete applications examples, such as AMR for gas dynamics, AMR incompressible flow solvers, and solvers for a variety of coupled fluid–particle problems (figure 6). The released software includes extensive design documentation and an automatically generated web-based reference manual using the doxygen system.

Applications

The APDEC Center has been involved with the development of applications in a variety of ways. Two projects—adaptive methods for low-Mach

number combustion and AMR for MHD models for tokamaks—were funded directly as part of APDEC. These projects have led to the development of substantial algorithmic capabilities as well as significant scientific investigations. APDEC is also collaborating with a variety of non-SciDAC investigators as part of the Center’s outreach activities (sidebar “Outreach Collaborations,” p24).

Simulating Flames

Turbulent combustion is a critical process for both energy and transportation. The detailed properties of turbulent chemistry interaction play a key role in both efficiency of the combustion process and emissions. In spite of its importance, high-fidelity modeling of turbulent combustion remains an elusive target. Detailed simulation of turbulent combustion processes involves the modeling of three-dimensional turbulent flow as well as the complexities associated with the thermochemical behavior of the system, such as reaction mechanisms, thermodynamic properties, and transport properties. These additional complexities not only make combustion simulation a challenging computational task, but also introduce a dependence on the existing experimentally determined characterization of the underlying physical processes.

In combustion research APDEC’s goal is to develop new predictive tools for modeling turbulent combustion processes. Specifically, the Center is modeling fluid–chemistry interactions with sufficient fidelity to predict not only the basic energetics, but also more detailed aspects of the behavior, such as the formation of pollutants. Simulation methods have been developed for solving the multispecies Navier–Stokes equations with chemical reactions that generalized an AMR algorithm for incompressible flow to the case of a low-Mach number hydrodynamics formulation to eliminate the time step constraint due to acoustic waves. The

The design of the APDEC AMR infrastructure has a strong focus on high performance.

Turbulent combustion is a critical process for both energy and transportation. The detailed properties of turbulent chemistry interaction play a key role in both efficiency of the combustion process and emissions.

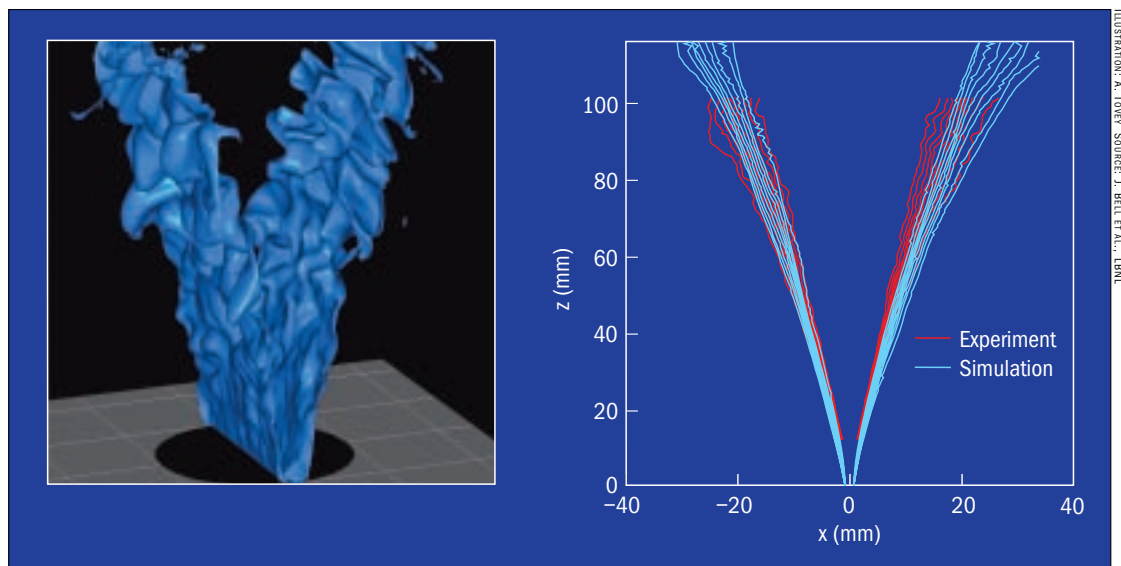


Figure 7. On the left, a flame surface image from the simulation of a laboratory-scale V-flame. On the right, a comparison between the numerical and experimental flame brush, which defines the envelope that confines the fluctuating flame.

resulting methods are able to resolve the range of length and time scales of laboratory-scale turbulent flames at a computational cost by approximately three orders of magnitude smaller than that of traditional direct numerical simulation approaches, based on explicit methods for uniform-grid compressible formulations. This simulation capability has been validated for a variety of two-dimensional laminar flames in studies of vortex-flame interaction and emissions from laminar diffusion flames. The methodology has also been used for the first three-dimensional direct numerical simulation of a premixed turbulent methane flame with detailed chemistry. More recently, scientists have simulated two laboratory-scale turbulent premixed flames—a rod-stabilized V-flame (figure 7), and a piloted slot Bunsen burner (figure 8). These detailed simulations captured the measured flame morphology as it evolved in the turbulent inlet flow, in terms of mean flame shape and velocity fields, global fuel consumption, flame surface density, and the shape and distribution of wrinkles in the flame surface (“Simulating Turbulent Flames,” *SciDAC Review*, Fall 2006, p25).

Magnetohydrodynamics

One of the central design issues for the next generation of magnetic fusion devices is the macroscopic stability of burning plasmas. In order to understand this problem, it is necessary to develop a collection of tools for fully nonlinear, time-dependent simulations. APDEC has developed AMR codes for ideal and non-ideal magnetohydrodynamics using a semi-implicit predictor-corrector approach. The explicit hyperbolic solver is a second-order Godunov method combined with an

implicit treatment of the diffusion terms. This method has been successfully applied to a variety of scientific problems. In pellet injection, the AMR single-fluid resistive MHD code successfully validated, albeit qualitatively, experimentally observed differences between high-field-side and low-field-side pellet launches (figure 9, p30). In magnetic reconnection, the instability of the current sheet and the formation and ejection plasmoids at high Lundquist numbers were observed with higher reconnection rates than those predicted by theory. Another scientific discovery was the suppression of the Richtmyer-Meshkov instability in MHD, which led to a new fundamental analytic theory to explain this phenomenon.

Improving Performance

AMR solvers for hyperbolic conservation laws already exhibit excellent scalability up to tens of thousands of processors. However, current implementations of iterative solvers for elliptic PDE on AMR grids scale only to 10^3 processors, due to the smaller amount of computation that is done between communication/synchronization steps. For the volume-of-fluid methods, there is the additional problem that the codimension-one irregular computations actually take up a great deal of time in the current implementation and lead to load imbalances that are not easily predicted. To deal with these problems, multiple rounds of performance improvement are being implemented, each consisting of the following three steps.

First, a suite of performance benchmark problems for a given capability is chosen, in consultation with the application users, and documented. Baseline measurements of serial

Higher-Order Finite-Volume Methods

It has long been known that, at least for simplified model problems, fourth/sixth-order accurate discretizations represent an optimum balance between accuracy and algorithmic complexity. Recent research developments in numerical methods indicate solutions to two of the major questions that have prevented the use of such higher-

order methods for multiphysics problems—the need for higher-order versions of finite-volume methods on locally-refined grids, and of semi-implicit temporal discretization methods. Similar ideas can be applied to obtain higher-order accurate embedded boundary methods for problems with complex geometries.




Figure 8. A flame surface from the simulation of a laboratory-scale slot Bunsen flame. The surface is colored by the local mean curvature.

performance and parallel performance for these benchmarks are made and documented. These benchmark problems are also incorporated into the APDEC performance regression test suite. Next, both serial and parallel code optimizations are made; in the process, new performance benchmark problems may be developed and incorporated into the performance test suite. At the end of this process, measurements of the optimized code are performed and documented again. In the last step, the optimized code is released for general use.

It is expected that each round of such performance improvements will lead to an order-of-magnitude improvement in the scalability of the APDEC solvers. In the initial round of performance improvements, the basic bulk-synchronous model of the parallelism and the associated APIs will be fixed. Within that framework, APDEC will introduce optimizations based on ideas such as

improvement in the locality of the distribution of grids onto processors, load balancing based on run-time measurements, and platform-dependent optimization of parameter choices, such as patch size and number of ghost cells. In later rounds it will probably be necessary to take more aggressive steps, such as modifying APIs to accommodate overlap of communication and computation, or taking advantage of specialized hardware features.

APDEC is also collaborating with other SciDAC Centers and Institutes (“SciDAC-2: The Next Phase of Discovery,” *SciDAC Review*, Spring 2007, p16) to provide access to tools specialized to support the discretization approach being used by APDEC. These include interfaces to solver packages from the SciDAC Towards Optimal Petascale Simulations (TOPS) Center, such as the *hypr* package for solving linear systems, and the Sundials package for solving ordinary differential equations (ODE) and differential algebraic equations (DAE) based on matrix-free Newton–Krylov methods. Additionally, visualization and data analysis tools from the SciDAC Visualization and Analytic Center for Enabling Technologies (VACET) are used for block-structured, locally-refined grids and embedded boundary discretizations of problems with complex geometries. Furthermore, storage resource management and high-performance input/output (I/O) libraries from the Scientific Data Management (SDM) Center (“From Data to Discovery,” *SciDAC Review*, Fall 2006, p28) are also very useful.

SciDAC Applications

Combustion

APDEC is continuing the development of combustion codes for laboratory-scale flames. The principal goals will be to couple the combustion codes to cold-flow simulations in complex geometries, and to begin to address the algorithmic problems associated with performing AMR calculations of low-Mach number reacting flows in closed containers. In the first problem, an embedded boundary incompressible Navier–Stokes solver for the

It is expected that each round of such performance improvements will lead to an order-of-magnitude improvement in the scalability of the APDEC solvers.

Low-Communication Poisson Solvers

Solutions to Poisson’s equation satisfy a strong form of elliptic local regularity. The local smoothness of the solution depends only on the local smoothness of the right-hand side with the coupling between disjoint regions mediated by real analytic functions with rapidly decaying derivatives. This suggests that an efficient parallel method for computing the potential induced by some charge distribution would be to compute local convolutions on a

disjoint union of patches, and then compute the smooth global coupling among the patches using a calculation with a much coarser, and less computationally expensive, discretization. While this observation has been used to implement fast particle methods with low communications costs, it had not been fully exploited for gridded discretizations. Traditional iterative methods for solving Poisson’s equation on locally refined grids have

unacceptably high communication costs, with communication/synchronization steps required after each local relaxation/operator evaluation step—that is, every few tens of floating point operations per grid point. However, recent developments in extending the potential-theoretic ideas used in the fast particle methods are leading to fast Poisson solvers with an order-of-magnitude increase in performance over traditional iterative methods.




Figure 9. A calculation of fuel pellet ablation in a plasma using the AMR MHD code. At the top are grids on full computational domain, and a blow-up of the refined region in the neighborhood of the pellet. Below, a comparison of high-field-side (top row) and low-field-side (bottom row) launch strategies is shown. The high-field-side injection exhibits greater penetration of the pellet into the center of the plasma, leading to greater fueling efficiencies.

Volume-of-Fluid Methods

One of the essential research developments driving the development of APDEC's embedded boundary methods is the systematic use of a formal truncation error and modified equation analysis as a basis for discretization design. The extension of this approach to moving and free boundaries will provide a solid

mathematical foundation for the application of volume-of-fluid methods to a broad range of sharp interface problems arising in reacting fluid flows (flame fronts), multiphase flow (surface tension, heat/mass transfer across phase boundaries), and conjugate heat transfer. For moving boundaries, researchers

are also developing robust and efficient second-order accurate level-set methods to represent the dynamics of sharp interfaces in fluids that can be combined with the existing grid generation methodology to compute sufficiently accurate finite-volume discretization information for moving fronts.

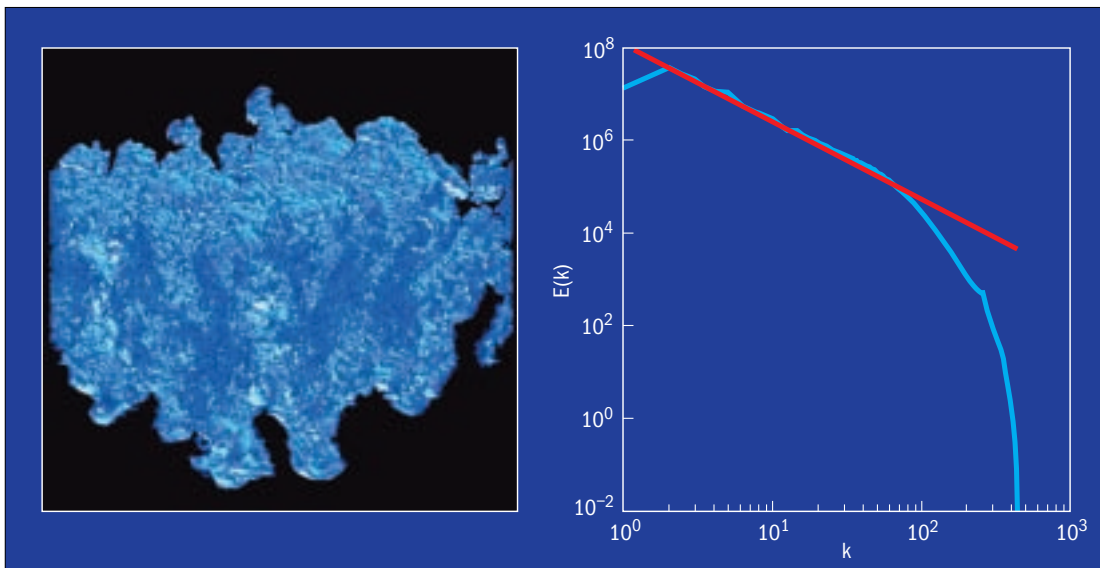


ILLUSTRATION: A. TOYER SOURCE: M. ZINGALES, SUNY-STONY BROOK

Figure 10. Image of a three-dimensional Rayleigh–Taylor unstable flame in a Type Ia supernova and the computed kinetic energy spectrum (blue curve) exhibiting the classical $k^{5/3}$ decay (red line).

flow inside the nozzle will provide the inflow boundary condition for the existing low-Mach number combustion solver in a rectangular domain. Coupling between the two solvers will be one-way with no feedback from the combustion solver on the cold-flow solution in the nozzle. The nozzle geometries will be those developed by experimental collaborators at LBNL, and validating the cold-flow simulations will be done using their experimental data. To address the issue of combustion in closed containers, APDEC researchers will implement a new low-Mach number reacting flow simulation capability in a closed domain based on a generalization of the allspeed projection formulation for fully compressible flow. Such an approach will allow users to correctly represent long-wavelength acoustics and their impact on combustion stability, while retaining the other advantages of the low-Mach number models. Finally, the latter capability will be extended to include an embedded boundary representation of engineering geometries to represent real engineering devices, such as gas turbines.

Magnetic Fusion

Researchers are continuing the development and extension of the AMR MHD code suite to simulate pellet injection and edge-localized modes (ELM) under APDEC. The main scientific goal is to achieve a predictive capability for the pellet injection phenomenon in ITER—the international fusion research project—including validation against pellet injection experiments. For the ELM simulations, the main scientific agenda is to simulate type I and type III ELMs and quantify the frequency of occurrence and the heat loads. Predicting these heat loads on the diverter is of clear practical importance for ITER. Another goal is to quantify the role of sharp pressure gradients and bootstrap current in the destabilization of MHD peeling and ballooning modes.

To support these goals, two approaches are being pursued. Initially, an existing semi-implicit AMR grid method is being used, while representing the vacuum region as a high-resistivity low-temperature plasma. Later, that approach will be extended to use a mapped grid discretization based on magnetic flux surface coordinates. The

Particle-in-Cell Methods for Electromagnetics

Particle-in-cell (PIC) methods are actually a hybrid of grid-based methods and a relatively singular representation of the interaction of the charged particles with the grid that take advantage of very specific features of uniform-grid discretizations of the fields to obtain stable and accurate results, rendering the extension to locally-refined grids problematic.

The standard methods used to solve the field equations preserve the divergence constraints of Maxwell's equations only on uniform staggered grids and with specific particle discretizations. The requirement that these properties be preserved in the presence

of locally refined meshes severely constrains the design choices for the discretizations. Another well-known problem regarding wave equations on locally refined grids is that outgoing waves on the fine grid that are not resolved on the next coarser level can get trapped on the fine grids. For particle methods, which are constantly generating such modes, the accumulated energy in those modes eventually overwhelms the solution, leading to large errors. New approaches to these problems include splitting the electric fields and currents into solenoidal and gradient components, and use the electrostatic PIC

algorithms described above to solve for the gradient component. Even though this will introduce elliptic solvers to a problem in which they previously did not appear, it is expected that, at least in some cases, the advantages of adaptivity will compensate for that. To solve the second problem, APDEC is taking advantage of the linearity of the problem in the development of new coarse-fine boundary conditions. These will permit outgoing waves on the fine grid that are not resolved on the coarse grid to propagate into a sponge layer surrounding the fine grid, where they will be harmlessly damped.

nonlocal electron heat flux will be modeled using a semi-analytical model, or by tabulating the electron heat flux from separate simulations of the Fokker-Planck equations.

APDEC scientists are also developing AMR MHD algorithms that treat the time scales associated with the fast magnetosonic and Alfvén waves implicitly, both based on semi-implicit formulations and using a fully implicit method based on the Newton-Krylov approach. In the latter case, the ideas can be used as the basis for designing preconditioners to remove the fast time scale dynamics.

Astrophysics

In the course of developing the AMR low-Mach number code turbulent combustion, the APDEC combustion team found that the same technology can be extended to treat the nuclear reaction networks and equation-of-state packages used to model nuclear burning in a Type Ia supernova ("Computing the Detonation of a White Dwarf Star," p10). The resulting code was used to study the microphysics of Rayleigh-Taylor unstable flames in two and three dimensions (figure 10). Following this work, APDEC will be collaborating with the Computational Astrophysics Consortium, led by Dr. Stan Woosley of the University of California-Santa Cruz, and the Adaptive Algorithms for Astrophysics Science Application Partnership (SAP), led by Dr. John Bell of LBNL. This project relies heavily on AMR as a simulation approach, and the SAP will support the development of new AMR capabilities not otherwise available as open source codes. These new capabilities will include: low-Mach number and allspeed models for fluid dynamics

with nuclear burning, AMR for multigroup flux-limited diffusion and other models for radiation, and AMR Poisson solvers for self-gravity. APDEC will work with these projects, particularly the SAP, to provide the infrastructure support for the development of these capabilities, particularly in the area of high-performance solvers for elliptic PDEs on AMR grids.

Accelerator Modeling

APDEC is developing a variety of tools for simulating plasma-wakefield accelerators. The embedded boundary compressible Navier-Stokes capabilities are used to simulate the gas dynamics inside of laser-irradiated capillary tubes. AMR capabilities for several models used to simulate the beam dynamics in a wakefield accelerator, such as an AMR particle-in-cell (PIC; sidebar "Particle-in-Cell Methods for Electromagnetics") algorithm for the QuickPIC approach, are also being developed.

Conclusions

The APDEC Center has been successful in providing software infrastructure for a broad range of applications involving PDEs. A key component of this success has been the mapping of mathematical concepts as they arise in the definition of models and algorithms to software components, leading to substantial reuse. APDEC is building on this model for success so that it can continue to develop and offer tools for enabling discovery in the future. ●

Contributor: This article was provided by APDEC Principal Investigator, Dr. Phil Colella of LBNL